



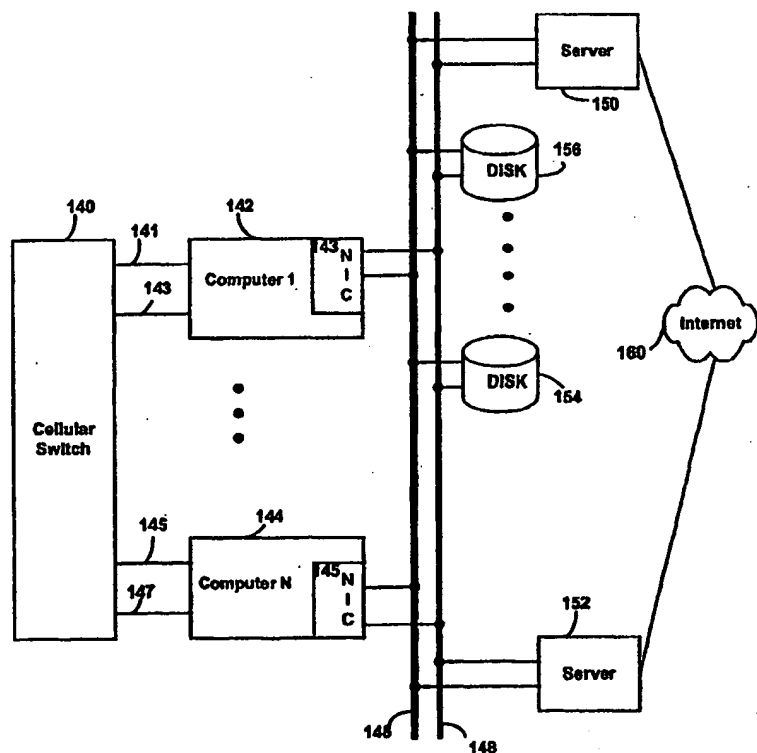
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H02H 3/05, H03K 19/003, H04B 1/74		A1	(11) International Publication Number: WO 00/22712
			(43) International Publication Date: 20 April 2000 (20.04.00)
(21) International Application Number: PCT/US99/23459		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 6 October 1999 (06.10.99)		<p>Published</p> <p><i>With international search report.</i></p> <p><i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	
(30) Priority Data: 09/169,838 9 October 1998 (09.10.98) US			
(71) Applicant: AT MOTION, INC. [US/US]; Suite 410, 303 Twin Dolphin Drive, Redwood Shores, CA 94086-1409 (US).			
(72) Inventors: EDMONDS, Paul; 2090 Bryant Street, Palo Alto, CA 94301 (US). ZHANG, Yi; 823 Groton Court, Sunnyvale, CA 94087 (US). XU, Chang; 1843 Poetry Drive, San Jose, CA 95131 (US). DOSHI, Priyen; 349 Kiely Boulevard #A306, San Jose, CA 95129 (US). CO, Stephen; 315 Andora Lane, San Ramon, CA 94583 (US). TEL, Michael, P.; 241 Arriba Drive #9, Sunnyvale, CA 94086 (US). CHAN, Andy; 157 Evergreen Way, Milpitas, CA 95035 (US).			
(74) Agent: TRAN, Bao, Q.; Fish & Richardson P.C., Suite 100, 2200 Sand Hill Road, Menlo Park, CA 94025-6936 (US).			

(54) Title: SERVER POOL FOR CLUSTERED SYSTEM

(57) Abstract

A computer system includes a plurality of interdependent processors for operating a common set of applications. Each interdependent processor executes an independent operating system image without sharing file system state information. Each interdependent processor has a network access card with a first network connection and a second network connection. The computer system includes a first active backplane (146) coupled to each first network connection of each processor; a second active backplane (148) coupled to each second network connection of each processor; the second active backplane (148) operating in lieu of the first active backplane (146) in case of a fail-over; and one or more directors coupled to the first and second active backplanes; each of the one or more directors load-balancing requests directed at a plurality of servers (150 and 152).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SERVER POOL FOR CLUSTERED SYSTEM

BACKGROUND

The invention relates to scalable and fault-tolerant computer systems.

5 The need for fast, reliable and secure access to vast amounts of shared data worldwide has been driving the growth of multiprocessing paradigm in which applications, data storage, processing power, and other resources are distributed among a pool of processors. A number of architectures have been developed over time to address the requirements of multiprocessing. Depending on the resources that
10 the processors share, multiprocessing architectures may be classified into three classes: share everything architecture, shared nothing architecture, and shared something architecture.

One example of a shared-everything architecture is a Symmetric Multiprocessing (SMP) architecture. An SMP system is capable of scaling
15 multi-process or multi-threaded loads so that application code can run on any processor in the system without software changes. Adding new throughput to the SMP system may be as simple as adding a new CPU board, provided the operating system can take advantage of it. Implementations of SMP generally provide a plurality of CPU and memory boards which communicate with each other and with
20 input/output boards over a wide and fast bus.

The SMP approach demands close communications between processors. The maintenance of consistency between processors is also non-trivial. The overhead and complexity of the consistency protocols may adversely affect the scalability of the SMP architecture. Further, processors in the SMP architecture typically share one
25 copy of the operating system. In addition to limiting the scalability of the SMP architecture, the sharing of the copy of the operating system creates a potential for many single points of failure occurring when many resources are shared.

One commonly used technique to provide fault-tolerance (fail-over) depends on a client application to recognize when a server is unavailable to satisfy a request,
30 and if so, either to locate another server or to deny the request altogether. For example, object techniques such as CORBA or Microsoft's Component Object Model

(COM) and Distributed Component Object Model (DCOM) may be used to implement this technique. These object architectures require a one-to-one relationship between the client and the server. If the server fails for any reason, the client needs to handle the failure either by finding another server that can perform the same service or by handling an error condition. These approaches require complex and time-consuming communication set-ups to provide sufficient fault tolerance for applications.

Another system supporting fail-overs for server processes, available from Microsoft Corp. of Redmond, Washington, is called Microsoft Cluster Server (MSCS). The MSCS system uses a hot-standby technique in which a primary server and a standby server send "keep alive" messages back and forth so that the standby server is activated if it cannot contact the primary server. This is a time consuming fail-over process. Further, the system is inefficient since computer resources of the standby server are not used until a failure occurs.

15

SUMMARY OF THE INVENTION

A computer system includes a plurality of interdependent processors for operating a common set of applications. Each interdependent processor executes an independent operating system image without sharing file system state information. Each interdependent processor has a network access card with a first network connection and a second network connection. The computer system includes a first active backplane coupled to each first network connection of each processor; a second active backplane coupled to each second network connection of each processor, the second active backplane operating in lieu of the first active backplane in case of a fail-over; and one or more directors coupled to the first and second active backplanes, each of the one or more directors load-balancing requests directed at a plurality of servers.

Implementations of the invention include the following. Each active backplane may be a switch. The switch may be an Ethernet switch. One or more networked data storage devices may be connected to the first and the second active backplanes. Further, one or more servers may be connected to the first or the second active backplane. Each director may be connected to each of the first and second

30

active backplanes. Each director may also be connected to a router. A peripheral device with an address may be coupled to the first or second active backplane. The address may be an Internet Protocol (IP) address. Further, the peripheral device may respond to the IP address when accessed from the first or second active backplane.

5 The address may also be a Media Access Protocol (MAC) address.

In a second aspect, a method for operating a computer system includes: executing an independent operating system image without sharing file system state information by each processor in a group of interdependent processors, each interdependent processor having a network access card with a first network
10 connection and a second network connection; transferring data on either a first active backplane coupled to each first network connection of each processor or a second active backplane coupled to each second network connection of each processor, the second active backplane operating in lieu of the first active backplane in case of a fail-over; and load-balancing requests directed at a plurality of servers using one or
15 more directors coupled to the first and second active backplanes.

Implementations of the method include the following. The transferring step includes routing data over each active backplane using a switch, which may be an Ethernet switch. Data may be accessed from one or more networked data storage devices connected to the first and the second active backplanes. Requests may be
20 communicated from one or more servers over the first or the second active backplane. Each director may be connected to each of the first and second active backplanes to provide load-balancing. Each director may also be connected to a router. A peripheral device connected to the first or second active backplane may be accessed at a predetermined address. The address may be a predetermined Internet Protocol (IP)
25 address, and the peripheral device may be accessed at the predetermined IP address from the first or second active backplane. The address may also be a predetermined Media Access Protocol (MAC) address.

Advantages of the invention include the following. The invention provides scalability and fault tolerance. The invention allows many servers to perform the
30 same task in an active/active scalable manner. The invention also supports load balancing among a pool of like servers. By providing a client process with access to a pool of like servers which are load balanced, the invention keeps the response time

for each request to a minimum. Thus, the invention supports high data availability, fast access to shared data, and low administrative costs through data consolidation. Additionally, the invention may be built using standard off-the-shelf components to reduce overall system cost.

5

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram illustrating a first computer system.

Figure 2 is a flowchart illustrating a process for accessing files.

Figure 3 is a block diagram illustrating a second computer system.

10 Figure 4 is a block diagram illustrating a computer node of the computer system of the present invention.

Figure 5 is a block diagram illustrating a third computer system.

Figure 6 is a block diagram illustrating a fourth computer system.

15 Figure 7 is a diagram of software supported by the computer system of the present invention.

Figure 8 is a diagram of a telephony module supported by the computer system of the present invention.

Figure 9 is a diagram of an internal services module supported by the computer system of the present invention.

20 Figure 10 is a diagram of an external services module supported by the computer system of the present invention.

DESCRIPTION

Fig. 1 shows a computer system in accordance with the present invention.

25 Telephone calls from a plain old telephone service (POTS) network are received by a cellular switch 100. The cellular switch 100 transfers the telephone calls to a plurality of computers 102 and 104 over multiple T1 line pairs 101-103 and 105-107.

The interface to the cellular switch 100 at each of T1 line pairs 101-103 or 105-107 receives analog voice signal from the network, bandpass-filters and
30 conditions the signal by the line interface and then applies the signal to a CODER/DECoder (CODEC) circuit. The CODEC filters, samples, and digitizes the inbound analog audio signal and passes this digitized audio signal for subsequent

digital signal processing (DSP). For outbound data, the interface expands stored, compressed audio data for playback; adjusts the volume and rate of speed of playback upon application or user request; and generates tones such as DTMF, MF, or any application-defined general-purpose tone as needed.

5 Each of computers 102 and 104 communicates with each other over an active backplane 110. Additionally, multiple data storage devices 106 and 108 are also connected to the active backplane 110. The active backplane 110 may be an Ethernet switch such as a Cisco 2900 switch, available from Cisco Systems, Inc. of San Jose, California. The active backplane 110 optimizes the Ethernet network by dividing it
10 into multiple smaller, more efficient collision domains, or segments. Individual segments created by the active backplane 110 can consist of one computer, like a file server, a group of computers, or even the entire business network. The active backplane 110 provides point to point bi-directional communication among many nodes such that each node has a full bandwidth access to the data storage device 106
15 or 108. Each of data storage devices 106 and 108 is a network Redundant Array of Independent Disk (RAID) data storage system. Each network RAID data storage system 106 or 108 connects directly to the active backplane 110. The network RAID data storage system may be a RAID array available from Network Appliances, Inc. which attaches directly to an Ethernet backplane 110 such as a 10Mbit, a 100Mbit, or
20 a 1Gbit Ethernet backplane 110.

 The network RAID data storage devices 106 or 108 includes a redundant RAID controller and is fully fault tolerant. As such, the network RAID data storage device 106 or 108 is a collection of disks under hardware or software control such that a single drive failure does not bring the system of Fig. 1 down. The network
25 RAID data storage device 106 or 108 also guards against fan and power failures and provides remote failure notification.

 Each of data storage devices 106 and 108 may be a RAID-1 system, in which every disk has a mirror image of its data stored on another disk. Alternatively, the data storage devices 106 and 108 may be a RAID-2 or RAID-3 sub-system which
30 stripes user data across a group of data drives (typically four or eight drives per group). The data storage devices 106 and 108 may also be a RAID-4 or RAID-5 sub-system which stripes block (or sometimes groups of blocks) of data and stores the

data entirely on an individual disk. Additionally, the operating system running on each data storage device 106 or 108 is optimized for file operations and has a thread-based non-preemptive multitasking environment. Data requests are transmitted over the active backplane 110 in accordance with a file sharing protocol, which is a special set of rules for communicating between a requestor running on the computer 102 or 104 and the network RAID data storage system 106 or 108 so that client applications in a computer may read and write to files on the network RAID data storage system 106 or 108. Further, applications may request services from server programs in a computer network. Through the file sharing protocol, an application can access files at a remote server as well as other resources, including printers, mailslots, and named pipes, to read, create, and update files on the remote server. The protocol can also communicate with any server program that is set up to receive a client request.

The file sharing protocol includes the Common Internet File System (CIFS) for Microsoft-based systems or the Network File System (NFS) for Unix-based systems. Alternatively, the file sharing protocols may be the Server Message Block (SMB) protocol, which is used over the Internet on top of its TCP/IP protocol or on top of other network protocols such as IPX or NetBEUI. The file sharing protocol supported by the RAID data storage device 106 or 108 provides a locking facility which may be a file locking facility or a byte-range locking facility. The locking facility enhances data integrity for the file sharing environment of Figure 1. Locking can be used to coordinate concurrent access to a file by multiple applications and users. It can prevent concurrent readers and writers of shared data from reading "stale" data (i.e., data currently in the process of being updated by another application) and/or overwriting each others' updates.

In file locking, the application requests that all other applications/users be denied read access to the file while the application holds the file open. If another application should attempt to try to open the file for reading (or reading and writing) while the file is already open with this deny-mode, the attempt fails. Alternatively, byte-range locking may be used to restrict other applications' access to sections of an open file, usually while the holder of the byte-range lock is intending to read or write the locked section. In certain implementations, byte-range locks can be obtained only

on already open files.

For interprocessor communications, a messaging system is provided with a unified mailbox containing message pointer lists and reference counts. The CIFS file system lock and application real-time programming processes such as open, lock, update, close are provided for each application requiring access to the mailbox. Lock management and disk allocation strategy is handled by the CIFS file system running on the RAID controller 106 or 108. In this way, each application telephony server and electronic mail server can operate on the messaging system directly without the need of a single threaded file or database server. Thus, maximum scalability is supported.

Moreover, the CIFS protocol automatically frees locks if a network connection is lost after some programmable period of time. Any computer 102 or 104 in the system illustrated in Fig. 1 can fail regardless of the locks it has outstanding and not affect the operation of any other computer that wants to operate on the same data. Hence, each computer or node in the cluster can operate on the entire file system independent of the state or operation of other nodes.

Fig. 2 is a flowchart illustrating a process 120 for accessing files stored on a data storage device is shown. The process 120 locks files at appropriate times such that any application can operate on the entire file system as if it owned the entire data storage device 106 or 108. The process 120 initially determines if the requested file in the particular data storage device has been locked (step 122). The file may be locked when another application has already accessed the file. If so, the process 120 waits for a predetermined period (step 124) before looping back to step 122 to check the file. If the file is available, the process proceeds to lock the file on the disk (step 126) and performs the requested data processing operation (step 128). Upon completion of the operation, the process 120 unlocks the file on the disk (step 130) before exiting (step 132).

Fig. 3 shows a second computer system with redundant backplanes. As in the system of Fig. 1, voice data from the POTS network is presented to a cellular switch 140. The cellular switch 140 in turn distributes the incoming calls to a plurality of computers 142 and 144. The cellular switch 140 communicates with each of computers 142 and 144 over multiple T1 lines 141-143 and 145-147, respectively.

The computers 142 and 144 in turn communicate with each other and with other peripheral devices over a redundant backplane with a first active backplane 146 and a second active backplane 148. The first and second active backplanes 146 and 148 are connected together.

5 Each of the active backplanes 146 and 148 may be an Ethernet switch such as the Cisco 2900 switch. The backplanes 146 and 148 divide the Ethernet network into multiple smaller, more efficient collision domains, or segments. Each backplane 146 or 148 provides point-to-point bidirectional communication among computers 142-144 such that each computer has a full bandwidth access to the shared peripherals
10 such as data storage devices 154 and 156. Instead of transmitting information to every device on the network and creating unnecessary traffic, each backplane 146 or 148 determines the exact destination for the data and transmits it only to the port where that device is attached. Each backplane 146 or 148 can also handle multiple transmissions at one time. In the system illustrated in Fig. 3, when one computer
15 sends information destined for another device on the network, the backplane 146 or 148 determines exactly where the data needs to go and creates a temporary connection between the two computers. This connection is a private link that is established only for the time needed to complete the data transmission. Also, unlike shared Ethernet hubs, the backplane 146 or 148 allows multiple simultaneous
20 telephone calls to be handled, thus increasing overall network performance. The backplanes 146 and 148 also provide redundancy in bus access. The use of the redundant backplanes 146 and 148 thus provides scalability for both processor-to-processor communication (as in node-to-node communication) and for processor-to-I/O communication (as in the case of a network-attached RAID device). The
25 processor-to-memory communication bandwidth limitation is obviated by the fact that nodes in a cluster provide processor-to-memory access within the node and do not require LAN access.

Also connected to the active backplanes 146 and 148 are servers 150 and 152. The servers 150 and 152 in turn are connected to the Internet 160. Additionally,
30 multiple network data storage devices 154 and 156 are also connected to the active backplanes 146 and 148 to satisfy data requests from computers 142 and 144 or servers 150 and 152.

The dual backplanes 146 and 148 provide scalability and fault tolerance. Further, multiple servers may perform the same task in an active/active scalable manner. By providing a client process with access to a pool of like servers which are load balanced, the response time for each request is optimized.

5 Fig. 4 illustrates the computer 142 in more detail. The computer 142 has a processor 172 which is connected to a processor bus 174. The processor 172 communicates with a random memory access (RAM) 176 and a read-only memory (ROM) 178 over the CPU bus 174. Additionally, the processor 172 communicates with one or more peripheral devices through an I/O controller 180 which is connected
10 to the CPU bus 174. The I/O controller 180 in turn provides an I/O bus 172 for communication with a display device 184, a data storage device 186, a telephone interface card 187, and a network interface system 188.

 The telephone interface card 187 may be a D/21H or a D/41H voice board, available from Dialogic Corporation of Parsippany, NJ, and which may support
15 digital SS7, digital E&M, digital loop start, digital DID, analog E&M, analog loop start, and analog DID. Each of the two (D/21H) or four (D/41H) loop start interfaces receives voice and telephony signaling information from the telephone network. Generally, using the interface, analog signal from the cellular switch 100 or 140 is bandpass-filtered and conditioned by the line interface and then applied to a
20 CODER/DECoder (CODEC) circuit. The CODEC filters, samples, and digitizes the inbound analog signal and passes this digitized signal to a digital signal processor (DSP). The DSP then uses Automatic Gain Control to compensate for variations in the level of the incoming analog signal, applies an adaptive differential pulse code modulation (ADPCM) or pulse code modulation (PCM) algorithm to compress the
25 signal and save disk storage space. The DSP also detects the presence of tones — DTMF, MF, or an application defined single- or dual-frequency tone. Thus, inbound telephony signaling (ring detection and loop current detection) are conditioned by the line interface and routed over a control bus to the DSP. The DSP responds to these signals, informs the application of telephony signaling status, and instructs the line
30 interface to transmit outbound signaling (on-hook/off-hook) to the telephone network. When recording speech, the DSP can use different digitizing rates from 24 to 64 kilobits per second (Kb/s) as selected by the application for the best speech quality

and most efficient storage. The digitizing rate is selected on a channel-by-channel basis and can be changed each time a record or play function is initiated.

Outbound processing is the reverse of inbound processing. The DSP processed speech is transmitted to the host processor for storage purposes. When playing back a stored file, voice information is received from the host processor and sent to the DSP, which converts the file into digitized voice. The DSP sends the digitized voice to the CODEC to be converted into analog voice and then to the line interface for transmission to the telephone network. For outbound data, the DSP expands stored, compressed audio data for playback; adjusts the volume and rate of speed of playback upon application or user request; and generates tones — DTMF, MF, or any application-defined general-purpose tone.

The network interface system 188 has a plurality of network interface cards 189-190. The network interface system 188 may be a Netelligent 10/100 TX PCI Intel UTP Controller, available from Compaq Computers Corporation of Houston, Texas. The Netelligent 10/100 unit can operate at 10 or 100 megabits per second (Mb/s) and automatically negotiates to the highest common speed when connected to a network. It has a network fault tolerance feature which enables its Controller to keep an active link by automatically switching to a backup network interface if a network connection fails.

In Fig. 5, another computer system is shown. A plurality of computers 200 and 202 communicate over a dual bus with a first bus 204 and a second bus 206. Additionally, a first Ethernet switch 210 is connected to the first bus 204. The Ethernet switch 210 may be the Cisco 2900 switch. The Ethernet switch 210 in turn is connected to an Internet pool of Web servers 214. Servers 214 support Web content retrieval, email, database management, and system management. The Ethernet switch 210 is also connected to a first director 216 as well as a second director 220. The first director 216 in turn communicates with an Internet router 222. The Internet router 222 is also connected to a hub 226 and to the Internet 160.

Correspondingly, a second Ethernet switch 212 is connected to the second bus 206. The second Ethernet switch 212, which may be the Cisco 2900 switch, is also connected to an Internet Web server 218 and to the second director 220. The second director 220 in turn is connected to an Internet router 224 which is connected to the

Internet 160. The Internet router 224 is also connected to the hub 226.

Directors recognize a Universal Resource Locator (URL) or Internet Protocol (IP) address as being associated with a pool of servers. If a server becomes unavailable, the server request is simply put in a slow poll mode and server requests are not sent to it until it starts responding. Directors provide various load-balancing algorithms to even out the load among a pool of servers. These devices assure high availability and scalability. By using directors 216 and 220 within such a clustered system, provisioning active/active pools of servers can be provided using off the shelf components to assure scalable, load balanced, fault tolerant access of clients to all server resources.

Fig. 6 shows yet another embodiment of the invention. Incoming data from the POTS network is received by a cellular switch 240. The cellular switch 240 in turn provides POTS data to computers 242 and 244 over multiple T1 connections for each computer. Each of computers 242 and 244 in turn is connected to a dual backplane system which is made up of a first active backplane 246 and a second active backplane 248. A plurality of servers 270 and 280 are also connected to the first active backplane 246 and second active backplane 248. Similarly, a plurality of data storage devices 250 and 252 are connected to the first active backplane 246 and the second active backplane 248.

Additionally, multiple directors 254 and 260 are also connected to the first active backplane 246 and the second active backplane 248. Each of the directors 254 and 260 is connected to each of routers 256 and 262. The routers 256 and 262 communicate with each other over a hub 258. The routers 256 and 262 are also connected to the Internet 160.

The deployment of the dual backplanes 246 and 248 and directors 254 and 260 allows many servers to perform the same task in an active/active scalable manner. By providing a client process with access to a pool of like servers which are load balanced, the system minimizes the response time for each request.

Turning now to Fig. 7, an overview of processes executing on the computer systems described above are shown. These processes include a telephony process 300 which is a collection one or more telephony application modules and an automatic speech recognition (ASR) server module. The telephony group 300 communicates

with an internal services process 320 via client/server communications. The internal services process 320 is a collection of software such as a CIFS application module, an application manager module, a facsimile manager module, a pager and message delivery manager module, a resource manager module, and an operation and administration (OA&M) manager module. The external services process 340 is a collection of services which handle tasks such as Internet Message Access Protocol 4 (IMAP4) /Post Office Protocol 3 (POP3) electronic mail, Short Message System (SMS) server, Web agent server, network and system management module, personal information management/electronic mail synchronization server module, and Web messaging server module.

The electronic mail server sends and receives messages using IMAP4 and POP3 SMTP protocols. The electronic mail server uses the same file store methods as the telephony application servers, thus supporting the abstraction of a unified mailbox to the phone as well as to the desktop. The Web server implements the Web user interface for message access, service management, system management, PIM synchronization, and Internet data collection. Through Active-X and Java applets or scripts, unified messages can be accessed over the Internet using Web browsers such as Microsoft Internet Explorer and Netscape Navigator. The user can use a Web browser to access his or her calendar, address book, and to-do list. The user can also access the Web server from the Internet to add or change services and configure service parameters. With the help of forms and menus, the user may interact with configuration parameters. Further, limited access to service parameters is supported over the telephone user interface.

The Personal Information Manager module manages a user's address book, calendar, and to-do list. An address book may contain entries with fields for name, title, company address, business phone, mobile phone, home phone, fax, electronic mail, Web page address and notes. A calendar may keep track of appointments, with features such as subject/agenda, location, date/time, and appointment length. A to-do list keeps track of things that need to be done. The to-do list entry may provide fields for subject/agenda, priority, due date/time, and reminder notes. The system can page, call or send electronic mail with to-do list content. PIM synchronization over the Web is provided to make information on palm-top computers available to the user at

any time anywhere by telephone or Web client access. The Web server is the gateway to handle this synchronization.

The telephony module 300 is illustrated in more detail in Fig. 8. In the telephony module 300, multiple client application module 302, 304 and 306 communicate with one or more automatic speech recognition (ASR) client module 308, one of which is assigned per communication port. The ASR client module 308 in turn communicates with an ASR server 310 which handles the speech recognition task. The ASR server software may be the Nuance-6 software, available from Nuance Communications of Menlo Park, California. The ASR server 310 enables speech recognition applications that allow users to talk to computers as if they were speaking with human agents. The ASR server 310 may utilize linguistic and statistical models to interpret and understand natural human speech to provide high recognition accuracy across a range of devices, languages, applications and vocabularies. The scalable client/server architecture of the system handles call volumes and vocabularies of all sizes while providing high reliability. Additionally, each client module 302, 304 and 306 also communicates over a telephony protocol 312 which supports facsimile, text, automatic speech recognition and voice processing.

Telephony application servers use the same file store methods as the electronic mail server. Users send and receive voice and fax messages and convert the electronic mail and Web text to speech using the telephony application. Electronic mail is sent to and from an IMAP4 electronic mail server which uses the same message store as the telephony application. In this way, the user is presented with the abstraction of a unified mailbox and advanced Web services both from the phone as well as the desktop.

When a server for telephony applications 300 goes down, all the ports in that node are lost. Since no shared services are implemented in the application servers, no fail-over is required. All ongoing calls on that server are dropped. The bus 204 or 206 will route new calls to the remaining telephony application servers.

Fig. 9 shows the internal services module 320 in more detail. A file sharing protocol handler 322 processes requests from client applications so that they can read and write to files or request services from server programs in the computer network. In this case, the file sharing protocol handler is a CIFS handler. By using CIFS

locking strategies, system wide resources and queues can be processed locally and in parallel with peer servers such as resource managers, application managers, pager/outdialer managers, and fax managers. The lock technique improves parallelism and scalability while reducing load balancing and fail-over complexity.

5 The CIFS file protocol handler 322 implements requests from one or more application commands such as allocate and lock. The application manager 326 allocates one or more universal ports, eliminates redundant inbound/outbound requests, and allocates and deallocates universal ports to set up various conference bridges as necessary.

10 The application manager 326 also communicates with a fax delivery manager 324 and a pager delivery manager 328. The fax delivery manager 324 maintains fax queues and is aware of the availability of dedicated and/or dynamically allocated dial-out ports. The fax delivery manager 324 is also responsible for retrying fax delivery if necessary. The fax delivery manager 324 communicates with the CIFS software
15 322 information regarding the fax queue management as well as the fax retrieval process. The pager delivery manager 328 maintains the pager queues and is aware of the dedicated and/or dynamically allocated dial-out ports. It also manages the retry of the pager delivery process. The internal services module 320 also contains a resource manager 330 which allocates the available conference bridges. The internal service
20 module 320 also has a OA&M (Operations, Administration and Maintenance) manager 332 which supports mailbox, port and disk utilities.

 Fig. 10 shows the external services module 340 in detail. The external services module 340 contains an IMAP4/POP3 electronic mail server 342. The server 342 handles electronic mail consolidation and unifies the mailbox access to the Web.
25 The IMAP4/POP3 server 342 communicates using an LDAP (Lightweight Directory Access Protocol) to an X.500 directory server 344. LDAP is a standard capable of providing open access to directory services on the Internet, as well as integrating directories and providing a global directory service. X.500 is a standard produced by the ISO/ITU defining the protocols and information model for a global directory
30 service that is independent of computing application and network platform. The X.500 standard defines a specification for a rich, global, distributed directory based on hierarchically named information objects (directory entries) that users can browse

and search. X.500 uses a model of a set of Directory Servers (DSAs), each holding a portion of the global Directory Information Base (DIB). The DSAs co-operate to provide a directory service to user applications in a way which means these applications need not be aware of the location of the information they are accessing.

5 The IMAP4/POP3 server 342 and the X.500 directory server 344 in turn communicates over the Internet 319. Further, the IMAP/POP3 server 342 communicates with an SMS server 346. Additionally, the X.500 directory 344 communicates with a personal information management (PIM)-electronic mail synchronization system 348. The PIM-electronic mail synchronizer 348 in turn
10 communicates with the Internet 319 such that a user can synchronize his PIM and electronic mail over the Internet.

 Additionally, the external services module 340 contains a network management module 350 which acts as an interface to the CIFS message store system. The network management module 350 also communicates over the Internet
15 319. A system administrator operating the network management module is provided with secure tools to manage the system over the network. Server management, network management, resource management and mailbox management can be all accessed from the Web.

 A Web messaging module 358 is provided to act as a service provider for
20 various message content presentations. The Web messaging software 358 also handles user preference configurations stored in the profile for system management and system configuration. Further, the external services module 340 includes a Web agent server 352 which operates in conjunction with a phone agent 354 to process spoken requests from the user for searching or operating applications over the
25 Internet. The phone agent 354 in turn provides commands to a Web methods interpreter 356. The Web methods interpreter 356 translates the verbal request or command to an Internet search and accesses the Internet 319 appropriately. The user can access Internet data such as stock quotes and weather reports from a cellular phone.

30 The above processes, in conjunction with the computer systems described above, operate in a fault-tolerant and scalable manner. Further, a pool of servers may execute these processes in an active/active scalable manner. The directors provide

load balancing among the pool of like servers. By providing the above processes with access to the pool of like servers which are load balanced, the system provides fast response time for each request.

Each of the above described modules or programs can be implemented in a high level procedural or object-oriented programming language to operate in conjunction with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

Each such computer program can be stored on a storage medium or device (e.g., CD-ROM, hard disk or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the computer to perform the procedures described. The system also may be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner.


Although a computer system with one or more active backplanes have been disclosed, other suitable network architectures may be used. Further, while the invention has been shown and described with reference to particular embodiments, those skilled in the art will understand that the above and other changes in form and detail may be made without departing from the spirit and scope of the following claims.

WHAT IS CLAIMED IS:

1. A computer system, comprising:
a plurality of interdependent processors for operating a common set of
5 applications, each interdependent processor executing an independent operating
system image without sharing file system state information, each interdependent
processor having a network access card with a first network connection and a second
network connection;
a first active backplane coupled to each first network connection of each
10 processor;
a second active backplane coupled to each second network connection of each
processor, the second active backplane operating in lieu of the first active backplane
in case of a fail-over; and
one or more directors coupled to the first and second active backplanes, each
15 of the one or more directors load-balancing requests directed at a plurality of servers.
2. The computer system of claim 1, wherein each active back plane is a switch.
3. The computer system of claim 2, wherein the switch is an Ethernet switch.
- 20 4. The computer system of claim 1, further comprising one or more networked
data storage devices coupled to the first and the second active backplanes.
5. The computer system of claim 1, further comprising one or more servers
25 coupled to the first or the second active backplane.
6. The computer system of claim 1, wherein each director is connected to each of
the first and second active backplanes.
- 30 7. The computer system of claim 1, wherein each director is connected to a
router.

8. The computer system of claim 1, further comprising a peripheral device coupled to the first or second active backplane, the peripheral device having an address.
- 5 9. The computer system of claim 8, wherein the address is an Internet Protocol (IP) address and wherein the peripheral device responds to the IP address when accessed from the first or second active backplane.
- 10 10. The computer system of claim 8, wherein the address is a Media Access Protocol (MAC) address and wherein the peripheral device responds to the MAC address when accessed from the first or second active backplane.
11. A method for operating a computer system, comprising:
executing an independent operating system image without sharing file system
15 state information by each processor in a group of interdependent processors, each interdependent processor having a network access card with a first network connection and a second network connection;
transferring data on either a first active backplane coupled to each first network connection of each processor or a second active backplane coupled to each
20 second network connection of each processor, the second active backplane operating in lieu of the first active backplane in case of a fail-over; and
load-balancing requests directed at a plurality of servers using one or more directors coupled to the first and second active backplanes.
- 25 12. The method of claim 11, wherein the transferring step includes routing data over each active backplane using a switch.
13. The method of claim 12, wherein the switch is an Ethernet switch.
- 30 14. The method of claim 11, further comprising accessing data from one or more networked data storage devices coupled to the first and the second active backplanes.

15. The method of claim 11, further comprising communicating requests from one or more servers over the first or the second active backplane.
16. The method of claim 11, wherein each director is connected to each of the first
5 and second active backplanes to provide load-balancing.
17. The method of claim 11, wherein each director is connected to a router.
18. The method of claim 11, further comprising accessing a peripheral device
10 coupled to the first or second active backplane at a predetermined address.
19. The method of claim 18, wherein the address is a predetermined Internet Protocol (IP) address, further comprising accessing the peripheral device at the predetermined IP address from the first or second active backplane.
15
20. The computer system of claim 18, wherein the address is a predetermined Media Access Protocol (MAC) address, further comprising accessing the peripheral device at the predetermined MAC address from the first or second active backplane.

**Fig. 1**




Fig. 2





FIG. 3

**Fig. 4**




FIG. 5




FIG. 6

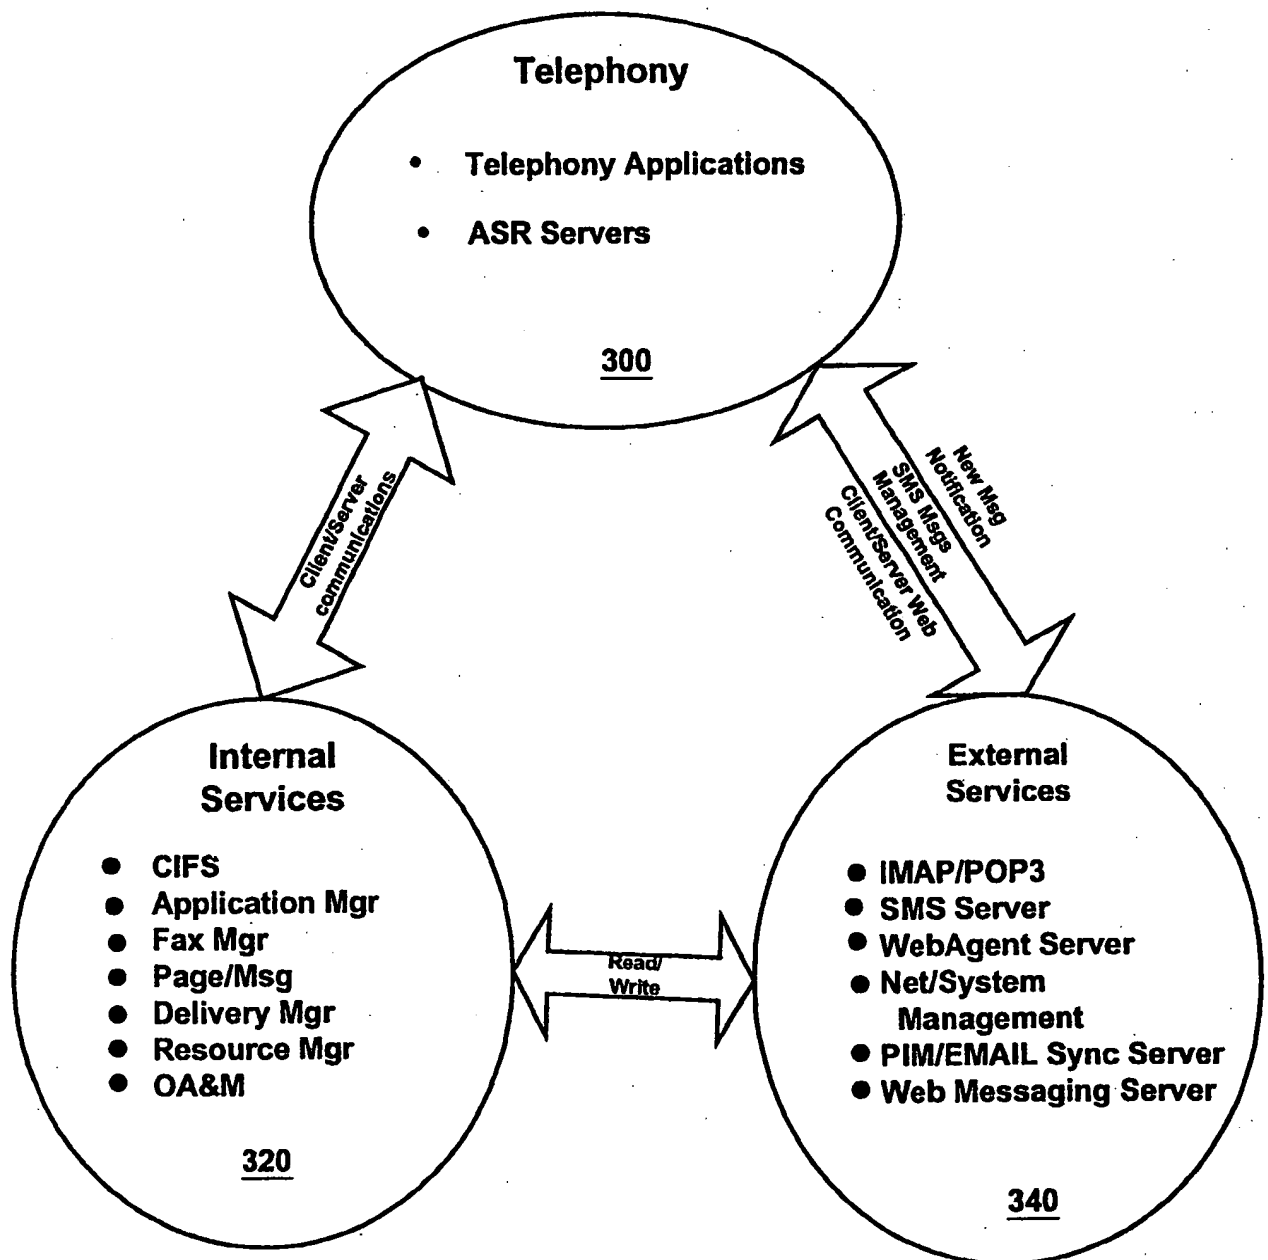


FIG. 7




FIG. 8

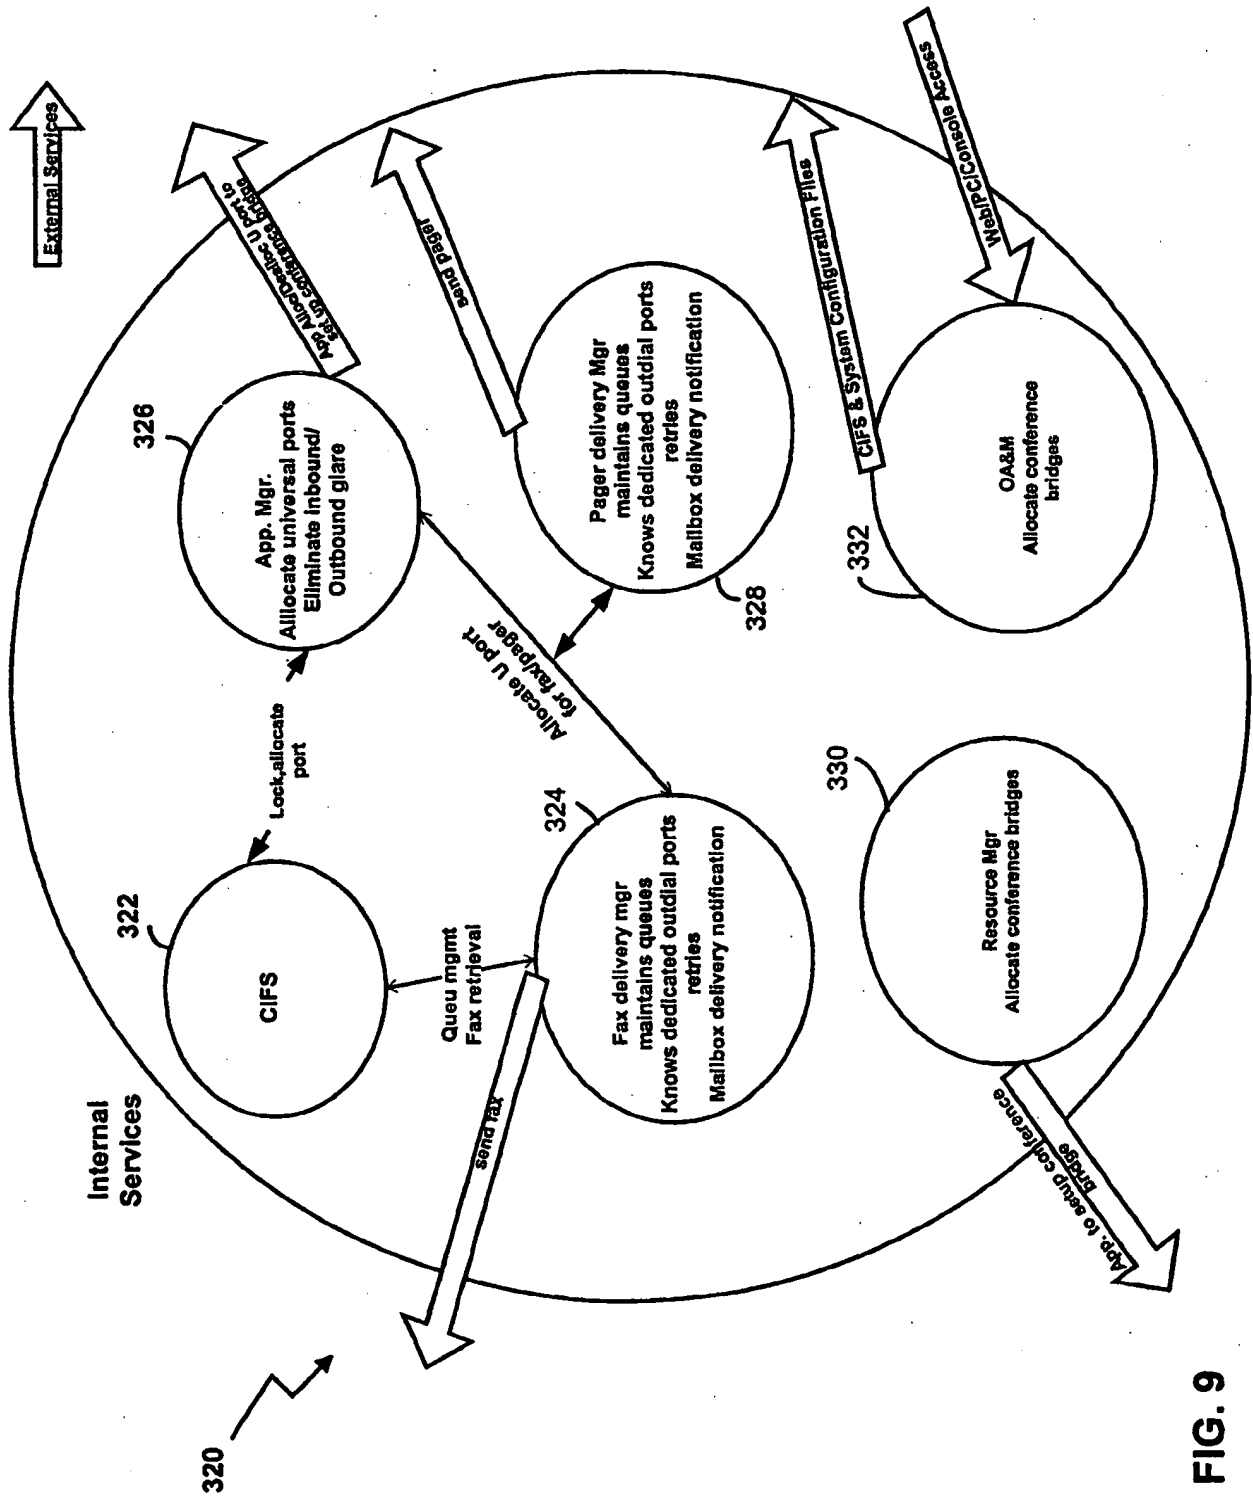


FIG. 9

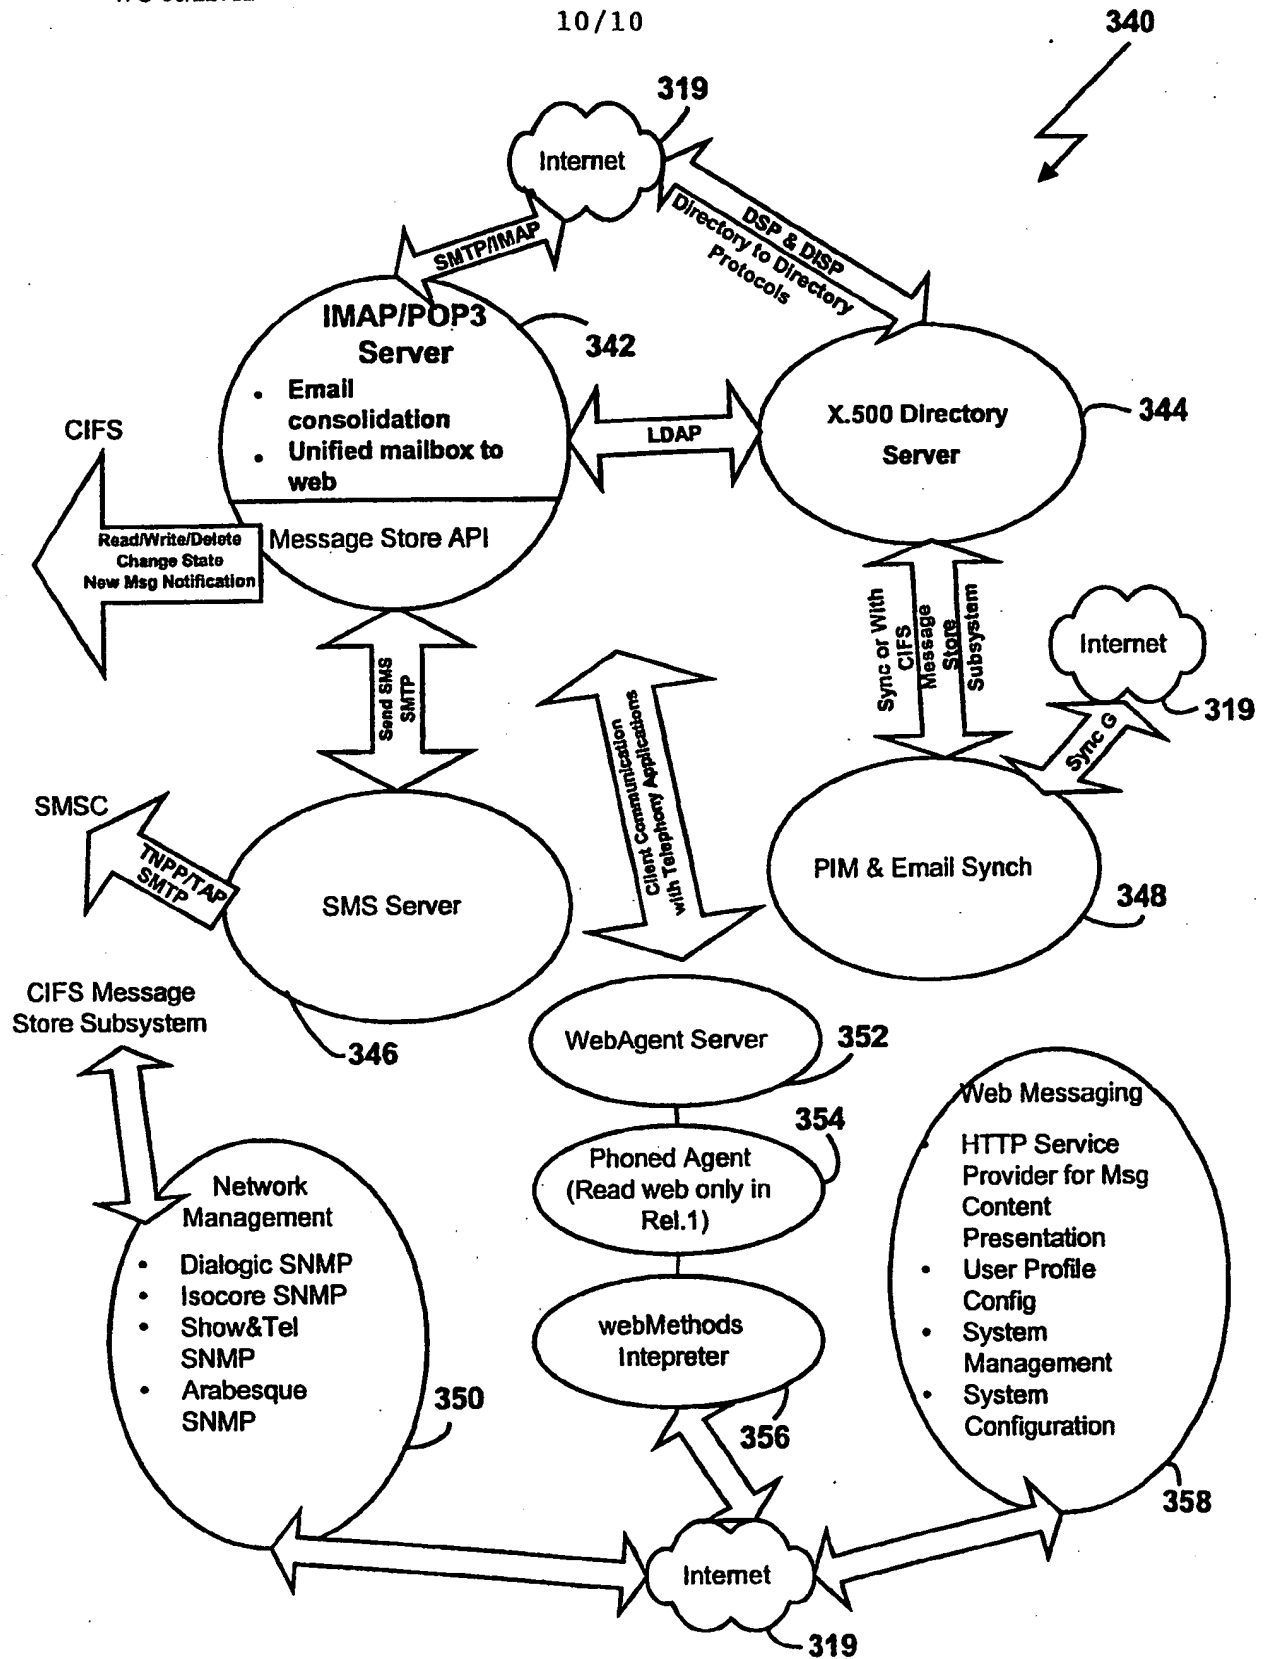


FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US99/23459

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H02H 3/05; H03K 19/003; H04B 1/74

US CL : 714/11, 12, 13, 6; 709/220, 249

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 714/11, 12, 13, 6; 709/220, 249

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
STN, EAST, WEST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 4,872,106 A [SLATER] 03 October 1989, see abstract, fig.2; col. 1, line 62 to col. 2, line 35.	1-20
Y	US 5,577,209 A [BOYLE et al.] 19 November 1996, see abstract, fig. 1; col.2, line 46-65.	1-20
Y	US 5,790,548 A [SISTANIZADEH et al.] 04 August 1998, see abstract, fig. 1; col.3, line 46 to col.4, line 43.	1-20
Y,P	US 5,852,724 A [GLENN, II et al.] 22 December 1998, see abstract, fig. 1; col. 1, line 49 to col. 2, line 9.	1-20
Y,P	US 5,935,230 A [PINAI et al.] 10 August 1999, see abstract, col. 5, lines 12-40.	1-20

☒ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

19 JANUARY 2000

Date of mailing of the international search report

07 MAR 2000

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

DIEU-MINH THAI LE

Telephone No. (703) 305-9408

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US99/23459

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US 5,996,086 A [DELANEY et al.] 30 NOVEMBER 1999, see abstract, fig.1, col. 2, lines 16-45.	1-20
A	US 5,768,623 A [JUDD et al.] 16 June 1998, see the whole patent.	1-20
A	US 5,805,785 A [DIAS et al.] 08 September 1998, see the whole patent.	1-20
A,E	US 5,968,185 A [BRESSOUD et al.] 19 October 1999, see the whole patent.	1-20

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.